

METHOD AND SYSTEM FOR
APPLICATION SPECIFIC PACKET FORWARDING

Claim for Priority

This application claims priority from European Patent Application No.
5 00120733.1, filed September 22, 2000, and which is hereby incorporated by reference as
if fully set forth herein.

Field of the Invention

The present invention relates to a method and system for forwarding data packets
through networks. More particularly, the present invention relates to a method and system
10 for handling information exchange through networks related to interactive broadcasting
systems, especially for interactive Internet based game shows.

Background of the Invention

A switched communication network transfers data from source to destination
through a series of network nodes. Switching can be done in one of two ways. In a circuit-
15 switched network, a dedicated connection is established through the network and is held
for as long as communication is necessary. An example of this type of network is the
traditional telephone system.

A packet-switched network, on the other hand, routes data in small pieces called
packets, each of which proceeds independently through the network. In a process called

store-and-forward, each packet is temporarily stored at each intermediate node, then forwarded when the next link becomes available. In a connection-oriented transmission scheme, each packet takes the same route through the network, and thus all packets usually arrive at the destination in the order in which they were sent. Conversely, each
5 packet may take a different path through the network in a connectionless or datagram scheme. Since datagrams may not arrive at the destination in the order in which they were sent, they are numbered so that the destination user can reorder them properly.

A data packet is routed through the network primarily according to its destination address. In order to determine the correct subsequent network the router has to convert
10 the destination address of a data packet into a corresponding next hop physical address (i.e. the outgoing port of a router). This task is called "address lookup" and is carried out as a part of the more complex "packet processing" task. The destination address is stored in a packet header. The packet header is a portion of a packet that precedes the actual data, containing source and destination addresses, error checking and other fields.

15 Packet processing , in addition, includes carrying out tasks like classification, filtering or load balancing, which may, based on multiple fields contained in the packet (not only the destination address), further influence the "address lookup" and the entire treatment and alterations applied to the packet in a router. For example, decide on specific QoS (Quality of Service) treatment of the packet, its mapping onto an MPLS
20 (Multiprotocol Label Switching) label, discarding it or sending it to a control point in case of filtering or splicing with another TCP (Transmission Control Protocol) connection in case of load balancing.

In the last decade the amount of data packet traffic being communicated over communication networks has grown exponentially especially the Internet a well-known

member of the family of connectionless packet-switched networks. In some circumstances the data packet traffic reaches such a level that conventional routers reach their limit. Since the performance of a router is crucial to the number of packets that can be transmitted through a communication network or from one communication network to another, a slow router can cause a backlog of data packets, or, in other words, a backlog of data packets is caused by an overload of the router. This might also happen because of an enormous amount of packets arriving at the router in a short period of time. Such a scenario might be produced if a huge number of users send messages over the network roughly at the same instant of time.

10 This might happen in various situations, for example, during a major sport event, when a large amount of visitors seeks to retrieve additional information about the event from a related web site. Furthermore, it has to be expected in systems providing interactive information exchange, such as Internet, television or radio game shows. In such shows the participants, which can practically be all viewers or listeners of the
15 respective show, are normally asked to respond to some action in the show, e.g., by answering a question, making a selection or by voting for something or somebody. Having a large number of participants, the communication system used for data communication might be overloaded because of the large number of responses being returned in a very short period of time.

20 A need thus exists for a method and a system that reduces the peak load of data transfer through the networks for interactive information exchange.

Summary of the Invention

5 The present invention provides a method and a system providing application specific forwarding of data packets through one or more networks. At least one server and a plurality of client machines are connected through the networks. The server is formed by a computer system which provides some service to the client machines, such as an interactive multimedia system. The expression 'interactive multimedia system' refers to any computer-delivered electronic system that allows the user to control, combine, and manipulate different types of media, such as text, sound, video, computer graphics, and animation. This also includes interactive broadcasting systems. Interactive broadcasting systems can, for example, be radio broadcasting, television broadcasting or cable television services with computer interfaces that enable viewers to interact with television programs.

15 The connection between client and server is normally by means of message passing, often over a network, whereby the present invention is not limited to any specific type of network. Practically, it can be a local area network (LAN), metropolitan area network (MAN), wide area network (WAN) with a wide variety of protocols, preferably using a store and forward protocol or a packet switching protocol, such as the Internet Protocol (IP).

20 First the server receives subscription messages from a subset of said client machines. The expression 'client machine' refers to the equipment a user or participant of the interactive multimedia or broadcasting system uses to interact with the transmitted program, such as a computer running a browser, i.e., a program which allows a person to read documents sent over the networks, or a set-top box, i.e., an electronic device designed to produce output on a conventional television set, on top of which it nominally sits. By receiving the subscription messages from the client machines the provider of the

particular program knows which users want to participate in the program. However, that does not necessarily mean that the provider actually knows the identity of the users, but only a temporary address specifying how to reach the client machine through the networks.

5 In response, the server composes a request message offering predetermined response options, whereby corresponding response messages are returned through the networks in one or more packets. In parallel, packet forwarding rules are set up in said networks specifying a particular treatment for said returned packets dependent on said predetermined response options. Finally, the server sends the request message to the
10 subset of client machines.

 In a preferred embodiment of the method and system according to the present invention the setting up of forwarding rules includes to configure one or more routers that forward said response packets. The configuring of the routers may include, individually or in combination, to instruct said routers to discard response messages containing
15 predetermined response options, to forward response messages containing a certain response option to a specified host connected to one of said networks, to combine more than one response messages arriving in a given time frame and to forward the combined messages as one message, to store the selected option of said response options in conjunction with the identity of the sender or to determine the amount of received
20 response messages for each response option.

 The method and system according to the present invention advantageously reduces the amount of packets sent through the networks, in order to transfer the information needed for the interactive broadcasting system. Thus, television stations or any other organizers of interactive broadcasting are enabled to use existing transmission

facilities that could not be used otherwise because of a limited transmission capability. So, the television stations could just lease such transmission facilities from ISPs (Internet Service Provider) for the short periods in which the show is on air. This is much more cost effective than using dedicated networks and servers. In other words, the transmission facilities do not need to provide a bandwidth capable to cope with the expected peak load, that would be unused for most of the time.

Furthermore, the method and system according to the present invention allows to use state of the art network facilities also for a nation or even world wide interactive program, such as a game show. The inventive features enable the transmission facilities to process the replies in 'real time', without having to tell the users to try again.

Most advantageously, the inventive method and system provides existing transmission facilities for coping with data traffic having a high peak load caused by potentially tens of millions of replies arriving at a server within the space of a number of seconds. In summary, one can say that, the taught method and system show a way of enabling current server technology to handle such a huge number of requests with real time constraints.

Brief Description of the Drawings

For a better understanding of the present invention, together with other and further features and advantages thereof, reference is made to the following description, taken in conjunction with accompanying drawings, and the scope of the invention that will be pointed out in the appended claims.

Fig. 1 shows a schematic view of a communication system to be used with the present invention;

Fig. 2 shows a high level block diagram of a router to be used with the present invention;

5 Fig. 3 shows a high level block diagram of a network processor to be used with the present invention.

Detailed Description of the Preferred Embodiments

With reference to Fig. 1, there is depicted a schematic view of a communication system 100 to be used with the present invention. The communication system 100
10 comprises a plurality of interconnected networks, whereas, for the sake of clarity, only a few networks are actually shown. A first, a second and a third local area network 102, 104 and 106 are shown in Fig. 1 each connected to a wide area network 110 through a first, a second and a third router 112, 114 and 116, respectively. Alternatively, the local area network may be formed by a domain, i.e., a group of computers whose hostnames
15 share a common suffix, the "domain name", wherein the hostname is a unique name by which a computer is known in a network.

The wide area network 110 itself might consist of more than one networks interconnected by a plurality of routers, as illustrated by a fourth router 118 that is in place of the plurality of routers. The first local area network 102 comprises a first, a
20 second and a third server 120, 122 and 124 all connected to the first router 112. For example, the first server is driven by the initiator of a game show or it belongs to the

domain of an ISP (Internet Service Provider) offering the environment for setting up interactive broadcasting.

The second local area network 104 includes a first, a second and a third client machine 130, 132 and 134 and the third local area network 106 comprises a fourth, a fifth, a sixth and a seventh client machine 140, 142, 144 and 146 belonging to users (not shown) of the communication system 100. A group or all of the users are participants in the game show.

The routers 112, 114, 116, 118 are entities that are capable of making forwarding decisions based on a bit pattern found in a packet to be used for transmitting information from one of the servers 120 to 124 to one of the client machines 130 to 146 and vice versa.

Now with reference to Fig. 2, there is depicted a high level block diagram of a router 200. The router 200 consists of a system processor 202, a switch fabric 204, a first, a second and a third adapter 206, 208 and 210 and a first, a second and a third line interface 212, 214 and 216.

The system processor 202 controls the operation of all remaining portions of the router 200. For the sake of clarity, the functional connections and communication lines to the remaining portions are not shown. The switch fabric 204 connects all adapters 206 to 210 to forward packets. The first, second and third adapter 206, 208 and 210 are connected to the first, second and third line interface 212, 214 and 216 respectively.

Each of the adapters 206 to 208 implement the functional behavior that is associated to Layer 2, 3 and 4 according to the OSI (Open Systems Interconnect) network

reference model. The OSI architecture is split between seven layers, from lowest to highest: physical layer (Layer 1), data link layer (Layer 2), network layer (Layer 3), transport layer (Layer 4), session layer (Layer 5), presentation layer (Layer 6), application layer (Layer 7). Whereas each line interface 212 to 216 is configured to provide Layer 1 functionality, i.e., the physical layer.

The Layer 1 activities perform line termination functions for links like Fast / Gigabit Ethernet, i.e., a local area network, or SONET/SDH (Synchronous Optical NETwork / Synchronous Digital Hierarchy), i.e., a broadband networking standard based on point-to-point optical fiber networks. The Layer 2 to 4 functions focus on the data link, network and transport functions, and are basically dealing with all activities required to forward a packet according to the packet's destination address and forwarding priority.

The three pairs of line interfaces and adapters 212 and 206, 214 and 208 as well as 216 and 210 form multiple parallel stacks of Layer 1 to Layer 4 functions, typically represented by a physical adapter card within the switch or router. As shown in Fig. 2, there are interconnected via the switch fabric 204. The system processor 202 provides general services for control and management of the router 200.

The router shown in Fig. 2 preferably deploys network processors for the packet forwarding functions in the adapters 206 to 210. In other words, the network processors performs the L2, L3 and L4 functions of the adapters 206 to 210 in Fig. 2.

Fig. 3 shows a high level block diagram of a network processor 300 to be used in a router according to Fig. 2. The network processor 300 comprises a general purpose CPU (Central Processing Unit) 302, routing and rules tables 304, a search engine 306, a

processor complex 308, a hardware (HW) assist 310, a scheduler 312, a packet buffer memory 314, a MAC (Media Access Control) frame processor 316 and an interface 318.

The function of each portion 302 to 318 of the network processor 300 and the overall operation will be apparent from a description of an example packet walk through.

- 5 A packet is received from the respective physical interface implementing the Layer 1 functionality. The MAC frame processor 316 implements the data link processing (Layer 2), extracts the Layer 3 PDU (Protocol Data Unit) - also called IP packet - and stores this IP packet in the packet buffer memory 314. The processor complex 308 consists of a single or multiple RISC-like (Reduced Instruction Set Computer) processor cores
- 10 optimized for packet processing. It fetches the network (Layer 3) and transport (Layer 4) packet headers from the packet buffer memory 314.

- The processor complex 308 coordinates all operations performed on a packet with the support from different coprocessors. The HW (Hardware) Assist 310 performs operations like checksum calculation. The search engine 306 performs the complex
- 15 address lookup function based on the routing and rules tables 304. In case of IP forwarding, the address resolution considers the IP destination address only, whereas, in case of Layer 4 switching, the IP destination and source address, and the TCP destination and source port numbers have to be considered for a unique identification of an end-to-end application flow identification. However, in case the router 300 is used as a dedicated
- 20 gateway machine with special security precautions on it, in order to function as a so called 'firewall', the operations require multidimensional rules to be evaluated prior to forwarding (or discarding) a packet. Example for such a rule-based forwarding may be:

IF IP source address is between 105.25.6.8 and 105.25.255.255 OR IP source address is bigger than 306.0.0.0 THEN Discard packet.

IF IP destination address is between 145.5.76.2 AND TCP source port equals 800 THEN Forward packet to output port 6 in priority 1.

- 5 In addition to the packet header fields, the processor complex 308 may also consider packet payload information for the forwarding decision. Each network processor or general purpose CPU core within a network processor is capable of dynamically updating routing or rule tables.

Moreover, it should be noted that the network processors' capabilities are not
10 restricted to pattern matching operations only. In principle, the programming model of network processors allows any algorithm to be executed for packet analysis or packet manipulation. However, this is only limited by feasibility of specifying the respective algorithm in a suitable number of elementary processor commands in order to fulfill given real time constraints.

15 The present invention will now be explained in greater detail, by way of an example, with reference to game shows. Game shows are ubiquitous in television schedules. Such games are characterized by a small set of members of the public physically present in an television studio replying to questions set by a presenter. The method and system according to the present invention allows every viewer to be a
20 contestant as well. Multiple choice questions could be given ("is it A, B or C ?") or voice recognition systems used to allow open questions ("What is the capital of France ?") and answers returned across a modem attached to a set top box attached to the television. This

enables viewers to become more active and in consequence less prone to change channels.

During the game, questions appear on the viewers screen and the viewers choose one. Different replies result in packets with different packet headers, or packets with
5 identical header information but different payload content. The actions to be taken by the router (or network processor) to support the Internet game application may be encoded in the regular routing or rules tables. For example, correct answers can be identified as a particular destination IP address and TCP port number pair. Correct answers are forwarded to a particular target host (the TV station) which keeps track of the scores of
10 the individual players while false responses are discarded. Alternatively, it is searched within the packet payload for a specific pattern to determine the correctness of the reply. All of these actions are provided by configuring one or more routers in the networks.

As the game show is bounded in time, and a winner must be found within that time, the order in which questions are asked can be altered as a function of their level of
15 difficulty and the number of participants left. The required reply handling rate for Internet based game shows according to the present invention is achieved by taking advantage of application specific knowledge within the forwarding of the packets, for example, during the packet treatment within a network processor. Based on the packet analysis, the network processor compiles different response messages to the originator of
20 the packet. Thus, an offloading of the game server is achieved.

Furthermore, in another embodiment of the present invention the response message to the questions are dispatched to the routers through which end users are connected to television stations servers and have those edge routers (routers 114 and 116 in Figure 1) make application specific forwarding decisions. Another alternative could be

to concentrate the response message processing to the one edge router the ISP or game show provider is connected to (router 112 in Fig. 1).

Further preferred embodiments of a particular response message treatment are now given. A router drops incorrect answers. Only correct answers are passed to the server, which does not need to do any checking but immediately registers all received replies as correct. Preferably, the checking occurs as close as possible to the game user (edge router), in order to reduce the overall bandwidth usage of the network. The server is configured in a way to mark answers as incorrect if it has not received a response message from a given user.

Alternatively, the router forwards correct and incorrect response messages to different locations, e.g., correct response messages to servers 120 and 122, incorrect responses to server 124. The aforementioned advantage also applies to this embodiment: the servers do not need to check the reply and not all the replies are sent over the same link. This leads to a reduction of the bandwidth requirement on a given link. Furthermore, the router receiving a response message determines its validity, so the data itself does not have to be forwarded, this further reduces the needed bandwidth.

Furthermore, a router could aggregate multiple small packets in one single large packet all containing either correct or incorrect replies. Reducing both the amount of data to be sent, the required bandwidth and the amount of processing that the server has to carry out.

Instead of configuring all routers in the path from on client machine to the receiving server, only the ISP-connected edge router (router 112 in Fig. 1) is instructed. This has the advantage that only a single router is required to be updated.

The Internet game show may or may not, be coupled with a live TV show. It is also conceivable that a game application, hosted and moderated by the ISP, runs an interactive game show. Contestants participate over the Internet until a winner is determined and the next game can start. Billing for the game participation can be coupled
5 with electronic money for the transmission of response messages by the participants.

The present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer system - or other apparatus adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general purpose computer system with a computer program that,
10 when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods.

15 Computer program means any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form.

20 It is to be understood that the present invention, in accordance with at least one presently preferred embodiment, teaches a method and system for handling interactive information exchange through networks having a plurality of client machines, the method comprising the steps of: composing a request message offering predetermined response options, whereby corresponding response messages are returned through said networks in

one or more packets; setting up packet forwarding rules in said networks specifying a particular treatment for said returned packets dependent on said predetermined response options; sending said request message to said subset of client machines.

5 Together, these elements may be implemented on at least one general-purpose computer running suitable software programs. These may also be implemented on at least one Integrated Circuit or part of at least one Integrated Circuit. Thus, it is to be understood that the invention may be implemented in hardware, software, or a combination of both.

10 If not otherwise stated herein, it is to be assumed that all patents, patent applications, patent publications and other publications (including web-based publications) mentioned and cited herein are hereby fully incorporated by reference herein as if set forth in their entirety herein.

15 Although illustrative embodiments of the present invention have been described herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various other changes and modifications may be affected therein by one skilled in the art without departing from the scope or spirit of the invention.